ATTRIBUTED IMAGE MATCHING USING A MINIMUM REPRESENTATION SIZE CRITERION

NAGW-1333

By:

A.C. Sanderson
N.J. Foster

Department of Electrical, Computer and Systems Engineering
Department of Mechanical Engineering, Aeronautical Engineering & Mechanics
Rensselaer Polytechnic Institute
Troy, New York 12180-3590

September 1989

CIRSSE Document #37

ATTRIBUTED IMAGE MATCHING USING A MINIMUM REPRESENTATION SIZE CRITERION

Arthur C. Sanderson

Electrical, Computer, and Systems Engineering Department, Rensselaer Polytechnic Institute, Troy, NY 12180

Nigel J. Foster

Electrical and Computer Engineering Department, Carnegie Mellon University, Pittsburgh, PA 15213

Manuscript submitted to IEEE Transactions on Pattern Analysis and Machine Intelligence

August 26, 1989

ATTRIBUTED IMAGE MATCHING USING A MINIMUM REPRESENTATION SIZE CRITERION

ABSTRACT

The interpretation of complex sensory data is fundamental to a variety of applications domains, and the matching of stored model structures to observed data from one or more sensors is an important approach to this problem. The minimum representation criterion is a metric of the overall complexity of a model and facilitates the unsupervised identification of model structure as well as parameters. This paper describes the use of this approach to the problem of matching noisy gary-level images to attributed models. Using the minimum representation criterion, the match between gray-level image features and an attributed graph model incorporates a representation size measure for the modeled points, the data residuals, and the unmodeled points. This structural representation identifies correspondence between a subset of data points and a subset of model points in a manner which minimizes the complexity of the resulting model. The minimum representation matching algorthm described in this paper is polynomial in complexity, and exhibits robust matching performance on examples where less than 30% of the features are reliable. The minimum representation principle is extendible to related problems using three-dimensional models and multisensor data matching

Index Terms: Attributed graph, correspondence, image matching, image models, representation size, shape recognition.

ATTRIBUTED IMAGE MATCHING USING A MINIMUM REPRESENTATION SIZE CRITERION

Arthur C. Sanderson

Electrical, Computer, and Systems Engineering Department,
Rensselaer Polytechnic Institute, Troy, NY 12180

and

Nigel J. Foster

Electrical and Computer Engineering Department,
Carnegie Mellon University, Pittsburgh, PA 15213

1 INTRODUCTION

Matching of models to image features is a fundamental step in computer vision systems. Such matching may take place at different levels of these systems, from template matching of raw images to symbolic matching of relational models. In this paper, we address the problem of matching localized spatial features with arbitrary attribute sets to either idealized or learned models. In mathematical terms, we match spatial patterns of points, where each point has an associated attribute vector with quantitative and symbolic values. The minimal representation criterion used to achieve an acceptable match is a principal topic of this paper.

Image matching is difficult to achieve with sufficient generality, speed, and robustness to be useful in practical systems. Many proposed algorithms are highly dependent on a choice of particular features and model representation, and they often require interactive or heuristic methods to extract features. Adding generality to matching procedures has been difficult particularly because evaluation functions or match quality measures do not generalize well. Image matching is inherently complex from a computational point of view, since the number of possible matches in general grows exponentially with the number of features. Polynomial complexity is an important property of any practical approach.

Good image matching algorithms must be able to handle feature uncertainty including missing data, extra features, and noisy attributes. This requirement has been particularly difficult to achieve since most evaluation functions are not able to handle missing or extra data in a consistent non-heuristic fashion. The representation criterion presented in this paper is inherently normalized to match size and number of attributes and directly accomodates missing and extra data.

This paper describes the minimal representation criterion [1,2,3,4] as a basis for image transformation and correspondence matching. We specifically address the problem of two-dimensional rigid, attributed point sets with missing and extra points. The algorithms developed are polynomial in complexity and near-optimal for this criterion. Examples of performance on highly variable gray-level images are shown.

Section 2 of this paper defines the image matching problem. Section 3 presents the minimal representation criterion principles. Section 4 describes a usually optimal, polynomial time algorithm for image matching and transformation. Section 5 presents some examples of the matching procedure.

2 ATTRIBUTED IMAGE MATCHING

Image matching problems have been approached using a variety of different hypothesize-and-test techniques in which potential matches are hypothesized and tested against evaluation criteria. These methods include template correlation [5], statistical pattern recognition [5], parameterized geometric fitting [6], and many different relational structure methods such as graph morphisms [7], compatibility graphs [8,9], and weighted relational matching [10]. In addition, heuristic techniques [11], Hough transform techniques [12], and relaxation labelling techniques [13] have been proposed. These references indicate examples of the various approaches, and a more detailed comparative discussion of these algorithms is included in [4]. The approach described in this paper is basically a geometric fitting technique which maps point sets to geometric models using a new metric for evaluating match quality. The minimal representation metric does not depend on the specific form of geometric modelling and is extendible to more general relational structure models.

In this paper, we consider images of rigid objects which have undergone arbitrary translation, rotation, and scaling in a two-dimensional plane parallel to the image plane. Each input image of an object is represented as a set of features with attributes, and each object model is represented in a similar manner for a given view of the object. In practice, the input image feature representation is extracted from the raw image data using other computer vision algorithms. The corresponding object model representation may be derived from a purely geometric model or by learning from a series of observations of input images. In addition to translation, rotation, and scaling, the image feature representation will include distortion, noisy attributes, missing (hidden or occluded) features, and added features. The image matching problem requires identification of the correspondence match between features and an associated geometric transformation which 'aligns' the image with the object model. The existence of an arbitrary transformation and the contribution of distortion and noise require a search over possible choices using an evaluation criterion which is tolerant to these effects. In this paper, the minimal representation criterion is used for the selection of the best correspondence and transformation.

An input image data feature representation consists of the ordered pair $D = (F, A)$ where $F = \{f_i,\ i = 1, \ldots, N\}$ is the set of feature labels, and $A = \{a_{ij},\ i = 1, \ldots, N,\ j = 1, \ldots, N_a\}$ is the set of feature attributes.

Each feature may have multiple attributes, and the set of attributes may differ among features. However, every feature in an image is required to have *(x,y)* position attributes denoted by

$$a_{i1} = u_i = x\text{-position of } f_i,$$

$$a_{i2} = v_i = y\text{-position of } f_i.$$

Similarly, an object model feature representation consists of an order pair $R = (G, B)$ where $G = \{g_i,\ i = 1, \ldots, M\}$ is the set of model feature labels, and $B = \{b_{ij},\ i = 1, \ldots, M,\ j = 1, \ldots, M_b\}$ is the set of model feature attributes. In this case

$$b_{i1} = x_i = x\text{-position of } g_i,$$

$$b_{i2} = y_i = y\text{-position of } g_i.$$

The attributes represented by A and B may be of four types:

1. *Positional* - (*x,y*)-position (required of every feature),
2. *Numerical* - numerical measures such as length, angle, area, curvature, number of neighbors,
3. *Symbolic* - symbolic labels such as color, texture,
4. *Relational* - relation of a feature to other features such as connected-to, on-top-of.

This data structure considers attributes independently and facilitates the development of the representation criterion which is strictly cumulative with respect to the set of features. For the problems considered in this paper, relational attributes will not be used. For highly noisy data relational attributes are difficult to incorporate into matching, and for rigid objects they are less useful since relative position is maintained by the rigid transformation.

2.1 Correspondence

Given an object model R and an input image I with data feature representation D, a match between them is defined by a correspondence and a transformation. The correspondence maps the model features G to the data features F. The transformation is the set of parameters which defineS the translation, rotation, and scaling used to geometrically align the corresponding features. In this paper, we assume that all correspondence matches are one-to-one, that is, one model feature matches to only one image feature and vice-versa. This assumption may be generalized, but simplifies the search problem and provides solutions which are more easily interpreted.

The *size* of the correspondence match, $N_m \leq min\,[M, N]$,is the number of model features which have a correspondence to a designated data feature. Not all model features have matches, and there may be added features in the image as well. The correspondence itself is expressed by the set of indices: $C = \{c_i,\ i = 1, \ldots, M\}$, , where

c_i = index of the image feature, f_{c_i}, which corresponds to the indicated model feature, g_i, when a match occurs and,

= 0, when no match occurs,

and $1 \leq c_i \leq N$. A particular correspondence match may therefore be represented by the ordered pair (i,c_i).

2.2 Transformation

Given a correspondence (i,i) where the model feature g_i is at point (x_i,y_i) and the image data feature f_i is at point (u_i,v_i), then the match is completely defined by a transformation T which transforms $(u_i,v_i) \rightarrow (x_i',y_i')$. In general, this transformation is defined by four parameters, $T = (t_u,\ t_v,\ O,\ s)$, where

$(t_u,\ t_v)$ = translation,
O = rotation angle,
s = scaling magnitude.

Fig. (1) illustrates such a transformation from (u_i,v_i) to (x_i',y_i'). While the data point is matched to the model point, the transformed data point does not necessarily align perfectly with the model point. The transformation will be derived from an evaluation criterion over a set of distorted and noisy data points, and will align relative to that global measure.

3 MINIMAL REPRESENTATION CRITERION

The minimal representation criterion [1,2,3] was introduced as an approach to unsupervised signal and data analysis in which the complexity of the data representation is used as a criterion for the choice of model structures and model parameters. The approach incorporates elements which express the complexity of the modeling procedure, the model size, and the size of the data residuals. In contrast to traditional mean square error measures of model fit which do not permit discrimination among model structures, the minimal representation size explicitly incorporates model structure and represents the tradeoff between complexity of the model structure and the resulting error in predicting the data points. This approach was demonstrated for several classes of parametric statistical models including evaluation of the order of an autoregressive model and determination of the number of clusters in a multivariate data sample. In [2], these techniques were applied to the unsupervised analysis of biomedical signals which resulted in a system for the automatic modeling, segmentation, and symbolic representation of complex patterns associated with medical diagnostic decisions.

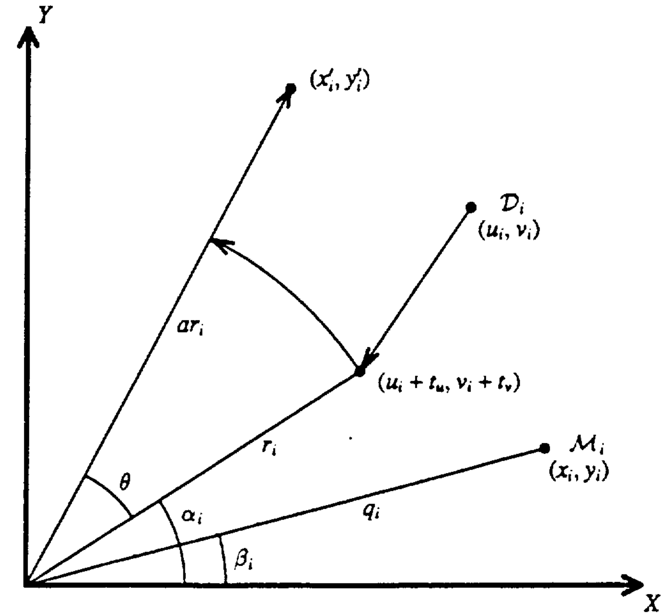

Figure 1 Transformation of an image point at (u_i,v_i) to a new point (x_i',y_i') using transformation T with four parameters: translation (t_u,t_v), rotation O, and scaling s.

The minimal representation criterion is based on a principle of minimum complexity of a program which explicitly regenerates observed data. Such a program includes a procedure, a model, and data residuals, and the size of the overall program is regarded as a measure of the complexity of the representation. In this approach, a simple model may require a complex data residual representation, while a more complex model will simplify the data residual representation. This tradeoff in overall complexity between model size and data residual size inherently provides a basis for choosing among alternative models. More generally, the procedure provides a tool for unsupervised decision-making.

Consider an observed data vector $\boldsymbol{x} = [x_1, x_2, \ldots x_N]$. The representation of this data vector is viewed as a program which regenerates the data points with some known resolution. In [1], this program is more formally defined in terms of a classic Turing machine model of computation. There may, in fact, be several different programs, π, which correctly generate the data points, and the 'correct' behavior of the system is regarded as the minimum size program p^* among these, such that

$$s(p*) = min\ s(p_i) \quad (bits)$$

where $s(\cdot)$ is the size of the program. As discussed in [1], the shortest program in an ensemble of such programs generated by a random process is the most likely program.

Each program, p, includes a number of segments which provide procedure code, model parameters, correspondence parameters, and data residuals. Each different algorithm or different model has a different set of program segments. In our previous work on clustering [1], for example, the model parameters included the cluster center positions in multivariate space, while the data residuals were encoded relative to these centers using a code which minimized the length of the data representation by encoding more probable (closer to the cluster center) data points with shorter length codes. In the image matching problem, the representation size $s(p)$ of each program includes the following terms:

$$s(p) = L + s(q) + s[C_q(\mathrm{x})] + s(e),$$

where

L = size of the program independent of the choice of model,

$s(q)$ = size of model parameters, including the transformation, the number of modeled data points, N_m, the correspondence match, and the feature attributes,

$s[C_q(x)]$ = [- $\log P_q(x)$], where

$C_q(x)$ = encoded residuals of modeled points, where

$P_q(x)$ = probability density function of the residuals of the modeled subset of observed data point attributes relative to the model q, and

$$s(e) = \sum_i \sum_j s(a_{ij})$$

is the representation size for the unmodeled data points. When all data points have uniform attribute sets, we can further simplify this to

$$s(e) = (N - N_m) \sum_j s(a_j), \\ = (N - N_m) S_a,$$

where S_a is the total representation size for the attributes of each unmodeled data point. In practice, S_a depends on the predefined resolution in bits of each of the attributes and is therefore usually fixed for a given problem.

The representation of the data residuals is based on an encoding which represents the more likely points by shorter code strings. There are many specific coding schemes which might be used, and we have implemented one such scheme which is based on a truncated hyperbolic distribution of errors. Incorporating this measure, we can write the representation size equation for a fixed model and data size in the following form:

$$s(p) = L \ + \ N_m log_2 M \ + \ \sum s_i \ + \ (N - N_m) S_a,$$

where

$$s_i = \sum log\left[w_{ij} E_{ij} \ + \ 1\right].$$

E_{ij} is the error due to the jth attribute at feature i,

$$E_{ij} = Error_j \left[g_i, f_{c_i}\right],$$

and w_{ij} is a weighting parameter which can be used to adjust the relative weight of attributes for different specific applications. For the image matching problem we have used Euclidean error measures as a basis for the encoding of position attributes, and the resulting representation size equation is

$$s(p) = L \ + \ N_m \, log_2 M \ + \ \sum log_2 \left[\left(\left(x_i' - x_i\right)^2 \ + \ \left(y_i' - y_i\right)^2\right)^{1/2} \ + \ 1\right] \ + \\ \sum_i \sum_{j=3}^{N_a} log_2 \left[w_{ij} E_{ij} \ + \ 1\right] \ + \ (N - N_m) S_a.$$

For the experiments described in this paper the second term $N_m \, log_2 \, M$ was considered a constant for each set of experiments.

4 IMAGE MATCHING ALGORITHM

For a given model and observed data, the best match is defined by an optimal transformation and an optimal correspondence between some subset of the data points and a subset of the model points. These two steps may be considered somewhat independently. An optimal transformation will exist for each possible correspondence which is chosen, and the algorithm must search over many possible correspondences in order to find an optimal match.

4.1 Transformation

The minimal representation transformation is in general quite different from the least mean square error transformation which is commonly used. A closed form analytical expression for the least mean squared error transformation may be derived and applied directly to a given model and subset of data points. The minimal representation match involves a logarithmic transformation of the square error terms and does not lend itself to a closed form analytical solution. We have used two algorithms for the calculation of the minimal representation transformation:

Numerical optimization - Partitioning of the search space using bounds on the volumes was implemented and combined with a random adaptive search for local minima. Hundreds of examples were studied using Monte Carlo techniques and the resulting transformations were examined and compared to mean squared error transformations. The minimal representation size results were stable and robust, particularly in the presence of added or missing data points.

Two-on-two transformations - It can be shown analytically that in one-dimension, a minimal representation transform always has two zero position error correspondences. In two dimensions, less than 1% of the optimal transformations found by simulation did not have two-on-two transforms, and for those cases the difference in the transformations was minor. We have therefore implemented a usually optimal transformation based on the two-on-two transformation. This approach dramatically decreases the complexity of the algorithm, reducing a continuous parameter search in four dimensions to a discrete search over $(N^2 - N)/2$ points.

4.2 Correspondence

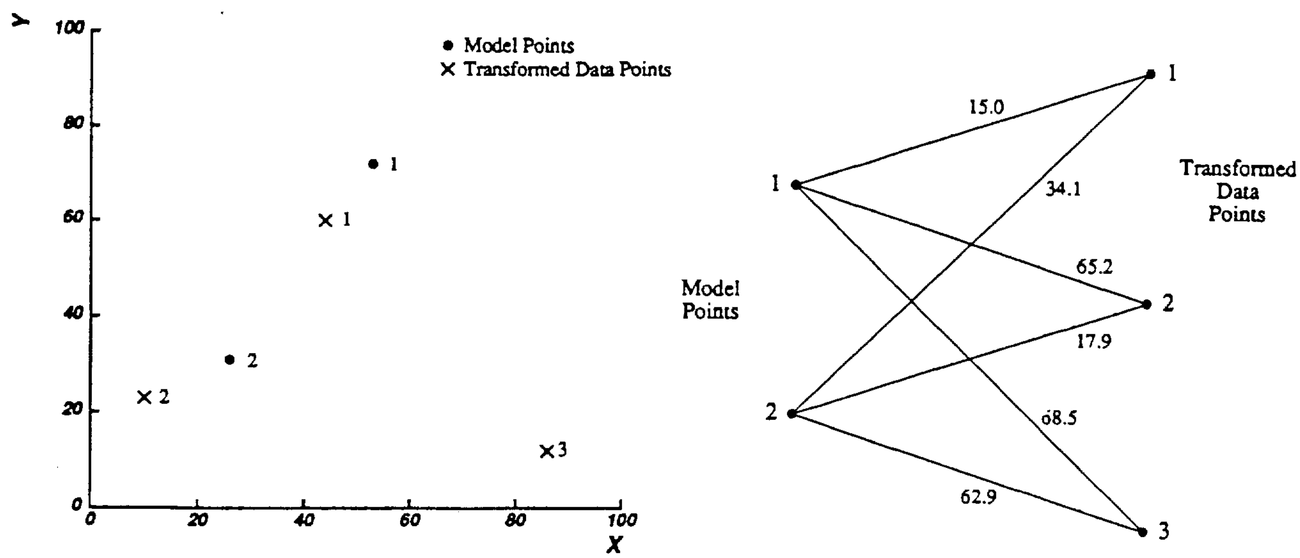

Figure 2 Two point sets and the bipartite graph of their possible correspondences.

The correspondence problem of finding the match between subsets of data points and model points which minimizes the representation size is solved by converting it to an assignment problem in the following form. Based on the minimal representation size equation, each pair of model and data points has two alternative representations. As a modeled point, the pair may have a representation size, S_p, associated with the model and residuals. As an unmodeled point, the pair will contribute a fixed size S_a. Fig (2) shows a set of model points, a set of transformed data points, and a graph of their possible interpoint mappings. The transformation parameters are not optimal and were chosen for the purpose of illustration. The point numbers do not indicate correspondence. The graph of interpoint distances is a complete bipartite graph, and the optimal correspondence can be viewed as an optimal assignment of left nodes to right nodes which minimizes the representation size.

In order to calculate the optimal correspondence, we

1. Assign the pairwise representation size to each arc of the complete bipartite graph.
2. Replace those representation sizes which are larger than S_a by the value S_a
3. Let $N' = max\ (M,N)$
4. If $M < N'$, add N'-M 'extra' nodes to the set of model nodes. Connect each extra model node to every data node using N arcs, each with weight S_a.

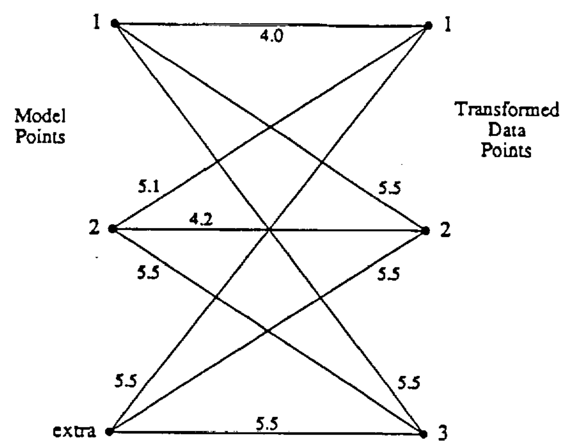

Figure 3 Expanded bipartite graph with representation sizes indicated as distance measures.

5. If $N < N'$, add N'-N 'extra' nodes to the set of data nodes. Connect each extra data node to every model node using M arcs, each with weight zero.

The resulting graph for Fig. (2) is shown in Fig. (3) with $S_a = 5.5$ bits. The optimal correspondence is now defined by choosing N' arcs such that (1) the sum of the arc weights is a minimum and (2) no two arcs share the same endpoint. A valid correspondence is indicated by a resulting arc weight which is less than S_a. All other arcs indicate that there is no correspondence between the two endpoints. The sum of the chosen arc weights is the representation size of the resulting match.

The assignment problem in a bipartite graph has been studied extensively [14], and a number of efficient algorithms exist. A straightforward solution would require evaluation of $N'!$ sets of arcs. Available algorithms typically are of order $O(N'^3)$ or $O(MN\ min(M,N))$ [15]. The latter algorithm was implemented here.

4.3 Complexity

The complexity of the resulting algorithm may be summarized as follows:

1. Compute optimal two-on-two transformations - $O(M^2N^2)$,
2. Compute the graph of representation sizes - $O(MN)$,
3. Compute the optimal match using the assignment algorithm - $O(MN\ min(M,N))$.

For large problems the computational complexity of the resulting algorithm is $O(M^3N^3\ min(M,N))$. While this algorithm still requires significant computation in its current form, on a typical size problem with $N = M = 30$, the computation is reduced relative to a brute force combinatorial algorithm by a factor of 1025. Many of the previous matching schemes have utilized heuristic techniques to reduce the computational complexity and did not optimize an objective measure of match quality. The algorithm described here produces usually optimal matches in polynomial time.

4.4 Improved Matching Efficiency

The performance of the basic matching algorithm can be improved using a number of algorithmic techniques and heuristics. The three methods summarized below utilize increasing assumptions about the characteristics of the data features.

1. **Precompute Representation Sizes:** The construction of the representation size graph requires the computation of *MD* representation sizes. Given a set of model points, it is possible to precompute all of the necessary representation sizes in a large *x-y* array. With such an array, the representation size calculation between the model point and any transformed data point is reduced to a single array access. Since a model is a collection of points, a number of separate arrays are required to represent all the possible representation sizes. The arrays are constant for a given model.
2. **Restrict Transform Space:** In most practical applications, there are fixed limits on the range of possible data point transformations. Those transforms which fall outside of this range can be ignored. In a typical vision application, the camera parameters are often fixed so that the scale of the data features is known within a few percent of their true value. With such scale, rotation, or translation restrictions, it is often not necessary to generate many of the candidate tranforms, and the search space is correspondingly reduced.
3. **Approximate Method:** In the basic matching algorithm, we explore all possible transforms without screening the candidate matches based on error criteria. This approach has provided an accurate view of the performance of the algorithm since it searches exhaustively over the candidates. In practice, one would like to reduce this search space based on prior screening of the errors. Such a fixed set of prescreened data points where only the most likely transformations and correspondences are explored, greatly reduces the search

problem and adapts it well to practical situations.

5 EXPERIMENTAL RESULTS

The matching algorithm described in this paper was tested on a variety of gray-level images with different degrees of complexity. Features extracted for matching are straight line segments and the vertices formed by the intersections and endpoints of such segments. The Popeye image processing system [16] was used to extract these edge-related figures by filtering, thinning, fitting of local line segments, logical reconnection, and simplification of the resulting line graph.

The line segments and their vertices are represented with a number of attached attributes. Each type of feature has a positional attribute. The position of a line segment is given by the center point of the line; while the position of a vertex is the point where two or more line segments intersect. In addition to the positional attribute, each segment also has a length attribute and a slope attribute. Vertex non-positional attributes include the number of line segments entering the vertex and the angle at which they enter.

Two examples of the feature extraction process are illustrated in Figs. (4) and (5). Fig. (4) shows a simple geometric shape with high contrast. The resulting edge-related features are clear and reliable as indicated by the dark lines and corner symbols in the figure. Fig. (5) shows a much more complex image which includes shading, highlights, and more subtle gray-tones. The resulting edge-related features are noisy and unreliable, and will often result in incomplete edge descriptions, or multiple vertices. Such complex images provide an important test of the minimal representation matching approach since they may contain a small percentage of repeatable features.

Fig. (6) shows an example of overlapping geometric shapes such as that in Fig. (4). These overlapping shapes provide a good test for the matching algorithm because they have occlusion among the objects. The contrast of the outer boundary of the shapes is still high, but the contrast among the objects is low and in general do not provide edge-related features. In these experiments, each of the shapes was matched independently of the others, so that no constraints among the group of objects were used. For the experiments, the independent shapes were matched with high reliability.

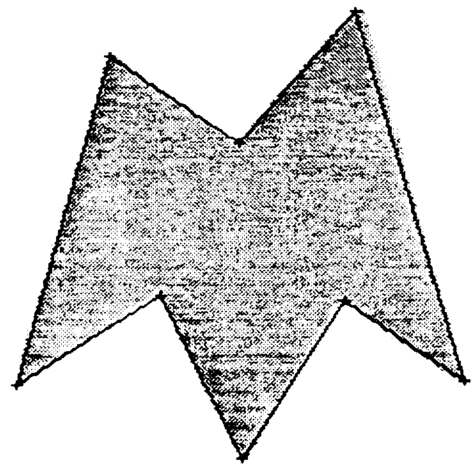

Figure 4 Feature extraction from a simple geometric shape with high contrast.

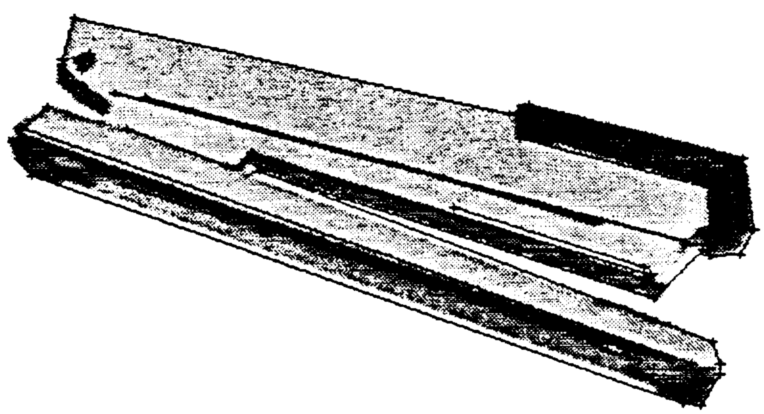

Figure 5 Features from a gray-level image of a three-dimensional object.

The effect of employing non-positional attributes was studied for these geometric shapes and the results of a study of images with simulated distortions is shown in Fig. (7). In each case, a random subset of features were selected from an image and a random set of synthetic features were added. Less than 50% of the features in all of these examples corresponded to the real image features. Four strategies were used on fifty examples of this type and the results are shown in Fig. (7). These results indicate that the algorithm is robust in spite of very large

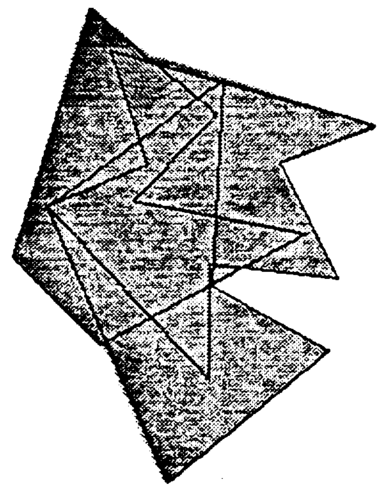

Figure 6 Example of minimal representation image matching with overlapping polygonal shapes. Models of the polygonal shapes were stored. Matching of each of the shapes to the gray level image was carried out independently.

Strategy	Vertex Attributes	Segement Attributes	% of Correct Matches
1	Position Only	Not Used	90.5
2	All	Not Used	99.5
3	Position Only	Position Only	96.0
4	All	All	100

Figure 7 Statistics for matching using several strategies with different incorporating different sets of attributes.

distortions of the data, and also that the addition of segment features, and the attributes for vertices and segments significantly improves the performance.

An example of a complex scene with an occluding object is shown in Fig. (8). The feature set derived from the original image is extremely noisy as shown in Fig. (8a). The correct match of the geometric model is shown in Fig. (8b).

Examples of matching to images of gray-level objects are shown in Figs.

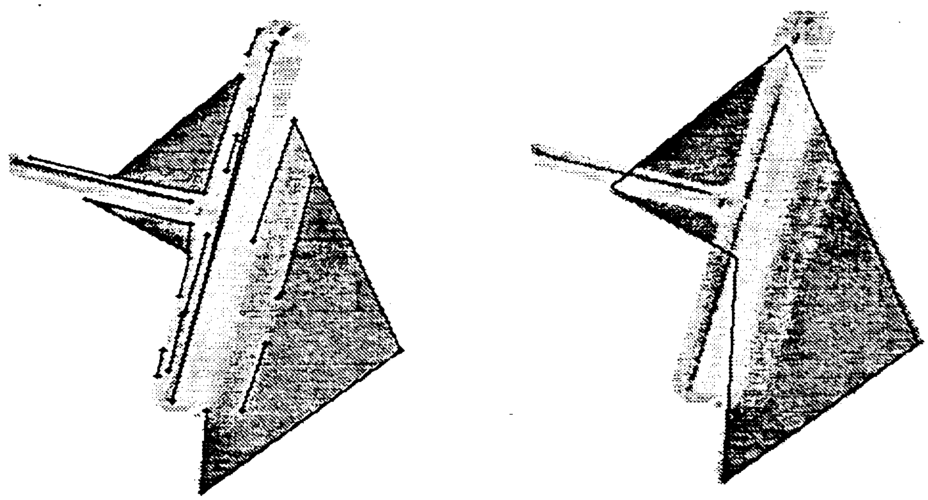

Figure 8 a. Example of image features for a noisy gray-level image of a polygonal shape with an occluding object. b. Correct matching of polygonal model to image features.

(9) and 10) for the example in Fig. (5). Fig. (5) shows the image and extracted features. Figure (9) shows the match of a model obtained from a slightly different angle of view. The resulting data image is quite noisy and varies significantly from the original model. The resulting match is still consistent with the model. Fig. (10) shows a match for an image of the object which is partially occluded. These noisy images typically had less than 40% consistent features as a basis for the match.

6 CONCLUSIONS

This paper has described a new approach to image matching which utilizes the minimal representation criterion as a means to obtain robust matching performance even when image data is extremely noisy. The results are encouraging in that they demonstrate consistent performance on samples of real gray-level images. The computational complexity of the approach is polynomial, but still large for applications such as inspection and robot control. Additional simplifications and approximations have been suggested which might make the technique feasible in these domains, and parallel implementation may be required to make the computation time acceptable.

The minimal representation approach to unsupervised decision-making is a general tool which has been employed in a number of different problems domains.

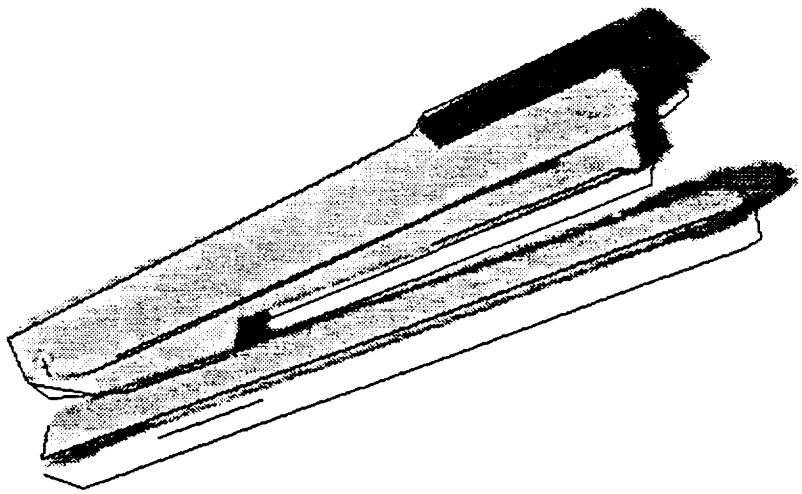

Figure 9 Matching of the gray-level image from Fig. (5) to a stored model obtained from the same object at a slightly different angle of view. Due to noise and distortion of the image, less than 30% of the features were consistent in this image, yet the algorithm was able to correctly match.

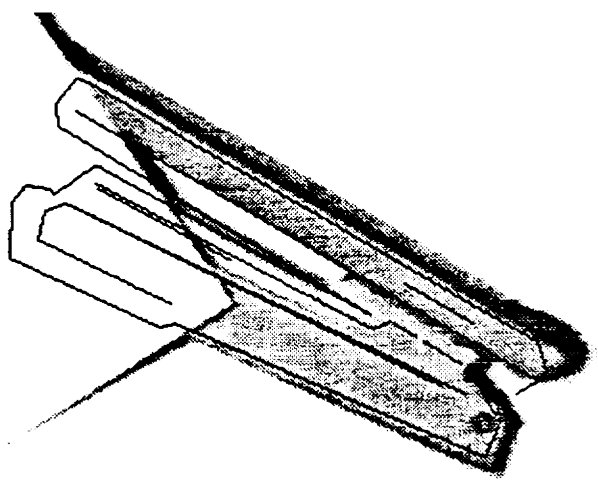

Figure 10 Matching of the stapler model to a gray-level image which is partially occluded.

The principle provides basic properties which seem to be useful in measuring and optimizing model structure as well as model parameters in a data interpretation framework. Such a minimal complexity or minimal entropy solution is appealing also from an intuitive point of view.

Extensions of this image matching approach to three dimensions, to problems with multiple sensors, and to problems with moving objects are all of particular interest for robotics applications. There are clear extensions to the work described here to all of these domains.

Acknowledgements

This work was supported in part by the Air Force Systems Command, Rome Air Development Center, Griffiss Air Force Base, New York 13441–5700 through the Northeast Artificial Intelligence Consortium (NAIC). Dr. Foster was supported in part by the Westinghouse Electric Corporation.

REFERENCES

[1] J. Segen and A. C. Sanderson, Model Inference and Pattern Discovery by Minimal Representation Method, Technical Report CMU-RI-TR-82-2, Carnegie Mellon University, Robotics Institute, Pittsburgh, PA 15213, July 1981.

[2] J. Segen, A. C. Sanderson, and E. Richey, "Hierarchial Modeling and Classification of EEG Signals," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-2, No. 5, 405-415, September, 1980.

[3] J. Segen, *A Pattern Directed Approach to Signal Analysis*, PhD Thesis, Carnegie Mellon University, 1980.

[4] N. Foster, *Attributed Image Matching Using a Minimum Representation Size Criterion*, PhD Thesis, Carnegie Mellon University, 1987.

[5] R. O. Duda and P. E. Hart, *Pattern Classification and Scene Analysis*, John Wiley and Sons, Inc., New York, 1973.

[6] H. S. Baird, *Model-Based Image Matching Using Location.* The MIT Press, Cambridge, Mass., 1985.

[7] M. A. Eshera and K. S. Fu. "A Graph Distance Measure for Image Analysis," *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-14(3):398-408, 1984.

[8] A. P. Ambler, H. G. Barrow, C. M Brown, R. M Burstall, and R. J. Popplestone, "A Versatile System for Computer-Controlled Assembly," *Artificial*

Intelligence, 6(2):129-156, Summer, 1975.

[9] R. C. Bolles and R. A. Cain, "Recognizing and Locating Partially Visible Objects: the Local Feature-Focus Method," *The International Journal of Robotics Research*, 1(3):57-82, Fall 1982.

[10] L. G. Shapiro and R. M. Haralick, "Structural Descriptions and Inexact Matchings," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-3(5):504-519, September 1981.

[11] N. Ayache and O. D. Faugeras, "HYPER: a New Approach for the Recognition and Positioning of Two-Dimensional Objects," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-8(1):44-54, January, 1986.

[12] J. Sklansky, " On the Hough Technique for Curve Detection," *IEEE Transactions on Computers*, C-27(10):923-926, October, 1978.

[13] L. Kitchen, " Relaxation Applied to Matching Quantitative Relational Structures," *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-10(2):96-101, February 1980.

[14] M. Gondran and M. Minoux, *Graphs and Algorithms*, John Wiley and Sons, Inc., New York, 1984.

[15] F. Bourgeois and J-C. Lassalle, "An Extension of the Munkres Algorithm for the Assignment Problem to Rectangular Matrices," *Communications of the Association for Computing Machinery*, 14(12):802-806, December, 1971.

[16] R. Bracho, J. Schlag, and A. C. Sanderson, Popeye: A Grey-level Vision System for Robotics Applications, Technical Report CMU-RI-TR-83-6, Carnegie Mellon University, Robotics Institute, Pittsburgh, PA 15213, May, 1983.